Asynchronous sampled-data synchronization with small communications delays

Gal Barkai   Leonid Mirkin   Daniel Zelazo

Technion - Israel Institute of technology

63rd IEEE Conference on Decision and Control

Networked multi-agent systems

- Multiple dynamic units interacting over a network
- Collective goals under limited communication resources
- Many applications
  - sensor networks
  - home automation
  - multi-robot coordination

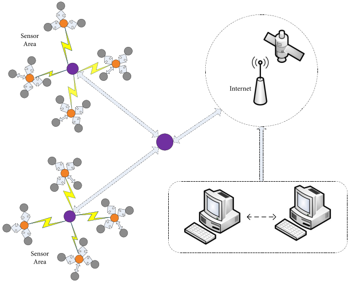

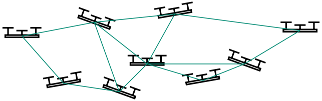

Archetype NMAS objective - State Agreement

Agents:

$$\dot{x}_i(t) = Ax_i + Bu_i(t), \quad i \in [1, \ldots, \nu]$$

with $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$.

Goal: design distributed control signals $u_i(t)$ such that the states synchronize, i.e.

$$\lim_{t\to\infty} \|x_i(t) - x_j(t)\| = 0 \quad \forall i, j,$$

with some trajectory generated by $\dot{r}(t) = A_0 r(t)$, for a given $A_0 \in \mathbb{R}^{n\times n}$ with $\mathrm{spec}(-A_0) \in \bar{\mathbb{C}}_0$, i.e.

$$\lim_{t\to\infty} \|x_i(t) - r(t)\| = 0 \quad \forall i.$$

- A_0 is an arbitrary model representing the target trajectory.

Archetype NMAS objective - State Agreement

Agents:

$$\dot{x}_i(t) = Ax_i + Bu_i(t), \quad i \in [1, \ldots, \nu]$$

with $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$.

Goal: design distributed control signals $u_i(t)$ such that the states synchronize, i.e.

$$\lim_{t\to\infty} \|x_i(t) - x_j(t)\| = 0 \quad \forall i, j,$$

with some trajectory generated by $\dot{r}(t) = A_0 r(t)$, for a given $A_0 \in \mathbb{R}^{n\times n}$ with $\mathrm{spec}(-A_0) \in \bar{\mathbb{C}}_0$, i.e.

$$\lim_{t\to\infty} \|x_i(t) - r(t)\| = 0 \quad \forall i.$$

- A_0 is an arbitrary model representing the target trajectory.

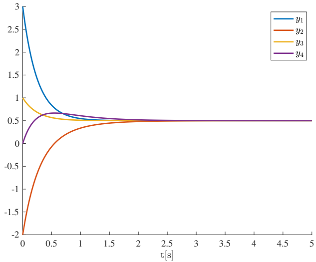

For $A_0 = A = 0$
$r(t) =$ const
Consensus problem

Archetype NMAS objective - State Agreement

Agents:

$$\dot{x}_i(t) = Ax_i + Bu_i(t), \quad i \in [1, \ldots, \nu]$$

with $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$.

Goal: design distributed control signals $u_i(t)$ such that the states synchronize, i.e.

$$\lim_{t \to \infty} \|x_i(t) - x_j(t)\| = 0 \quad \forall i, j,$$

with some trajectory generated by $\dot{r}(t) = A_0 r(t)$, for a given $A_0 \in \mathbb{R}^{n \times n}$ with $\mathrm{spec}(-A_0) \in \bar{\mathbb{C}}_0$, i.e.

$$\lim_{t \to \infty} \|x_i(t) - r(t)\| = 0 \quad \forall i.$$

- A_0 is an arbitrary model representing the target trajectory.

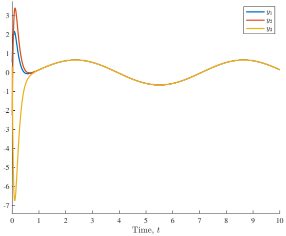

For $A_0 = A$
$r(t)$ is time-varying
Synchronization

Communication constraints

The controller must respect three types of constraints:

Communication constraints

The controller must respect three types of constraints:

Spatial constraints: an agent can transmit only within its neighborhood $\mathcal{N}_i(t)$. The neighborhoods can be time-varying and communication directed.

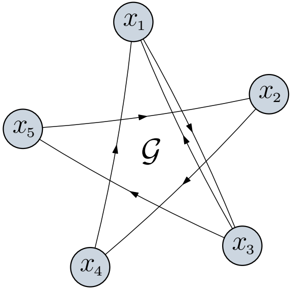

Communication constraints

The controller must respect three types of constraints:

Spatial constraints: an agent can transmit only within its neighborhood $\mathcal{N}_i(t)$. The neighborhoods can be time-varying and communication directed.

Temporal constraints: agents communicate only at discrete sampling instances, $t \in \{s_k\}$. The sampling can be aperiodic and asynchronous.

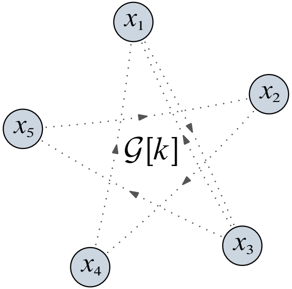

Communication constraints

The controller must respect three types of constraints:

Spatial constraints: an agent can transmit only within its neighborhood $\mathcal{N}_i(t)$. The neighborhoods can be time-varying and communication directed.

Temporal constraints: agents communicate only at discrete sampling instances, $t \in \{s_k\}$. The sampling can be aperiodic and asynchronous.

Transmission delays: the information transmitted by agent j at $t = s_k$ arrives to agent i at $t = t_{ij}[k] := s_k + \tau_{ij}[k]$. The delays can be time-varying and heterogeneous.

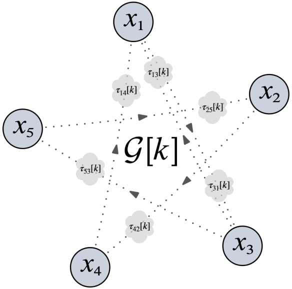

These pose significant challenges for control design.

Basic assumptions

- $\mathcal{A}_1$: The pair (A, B) is stabilizable and there exists a gain $\bar{F}$ such that $A_0 = A + B\bar{F}$.

Solvability

The agents are stabilizable and the trajectory is attainable.

Basic assumptions

- $\mathcal{A}_1$: The pair (A, B) is stabilizable and there exists a gain $\bar{F}$ such that $A_0 = A + B\bar{F}$.
- $\mathcal{A}_2$: there is a strictly increasing sub-sequence of sampling indices $\{k_p\}$ such that for all $p \in \mathbb{Z}_+$
  1. the intervals $s_{k_{p+1}} - s_{k_p}$ are uniformly bounded;
  2. $\bigcup_{k=k_p+1}^{k_{p+1}} \mathcal{G}[k]$ contains a directed rooted tree.

Joint connectivity

Information persistently propagates through the network

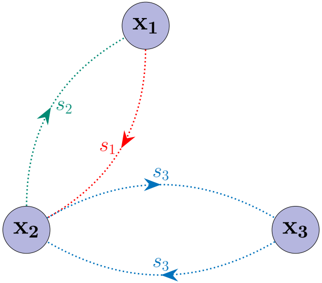

Basic assumptions

- $\mathcal{A}_1$: The pair (A, B) is stabilizable and there exists a gain $\bar{F}$ such that $A_0 = A + B\bar{F}$.
- $\mathcal{A}_2$: there is a strictly increasing sub-sequence of sampling indices $\{k_p\}$ such that for all $p \in \mathbb{Z}_+$
  1. the intervals $s_{k_{p+1}} - s_{k_p}$ are uniformly bounded;
  2. $\bigcup_{k=k_p+1}^{k_{p+1}} \mathcal{G}[k]$ contains a directed rooted tree.
- $\mathcal{A}_3$: incoming information is time stamped and

$$s_k + \tau_{ij}[k] < s_{k+1}, \quad \forall i, j \in \mathbb{N}_\nu,\ k \in \mathbb{Z}_+.$$

Small delays

The delays can be locally calculated and are *small* compared to the sampling interval.

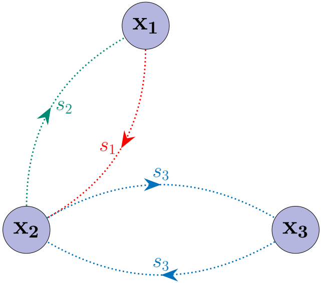

The standard approach to agreement problems

Sequential design:

1. Transmit the state, $x_i(t)$, and use a consensus-based structure to satisfy the spatial constraints.

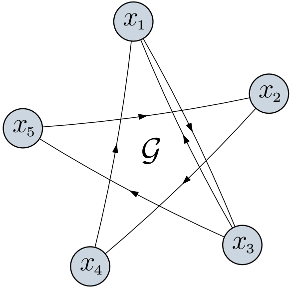

$$u_i(t) = K \sum_{j \in \mathcal{N}_i(t)} (x_i(t) - x_j(t))$$

The standard approach to agreement problems

Sequential design:

1. Transmit the state, $x_i(t)$, and use a consensus-based structure to satisfy the spatial constraints.
2. Assume the sampling is periodic and the control signal constant between samples.

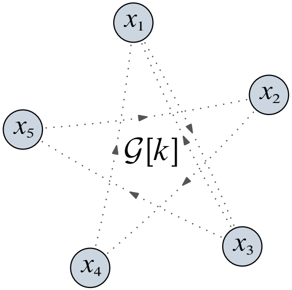

$$u_i(t) = K \sum_{j \in \mathcal{N}_i[k]} (x_i(s_k) - x_j(s_k))$$

The standard approach to agreement problems

Sequential design:

1. Transmit the state, $x_i(t)$, and use a consensus-based structure to satisfy the spatial constraints.
2. Assume the sampling is periodic and the control signal constant between samples.
3. Treat delays and sampling variations as perturbations.

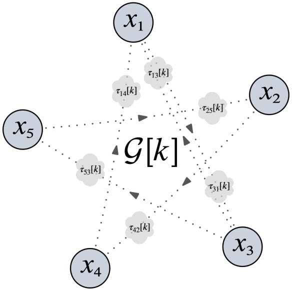

$$u_i(t) = K \sum_{j \in \mathcal{N}_i[k]} (x_i(s_k) - x_j(s_k))$$

The standard approach to agreement problems

Sequential design:

1. Transmit the state, $x_i(t)$, and use a consensus-based structure to satisfy the spatial constraints.
2. Assume the sampling is periodic and the control signal constant between samples.
3. Treat delays and sampling variations as perturbations.
4. Solve a robust control problem to find an appropriate gain.

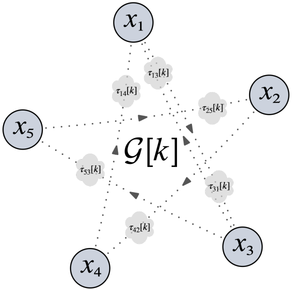

$$u_i(t) = K \sum_{j \in \mathcal{N}_i[k]} (x_i(s_k) - x_j(s_k))$$

The standard approach to agreement problems

Sequential design:

1. Transmit the state, $x_i(t)$, and use a consensus-based structure to satisfy the spatial constraints.
2. Assume the sampling is periodic and the control signal constant between samples.
3. Treat delays and sampling variations as perturbations.
4. Solve a robust control problem to find an appropriate gain.

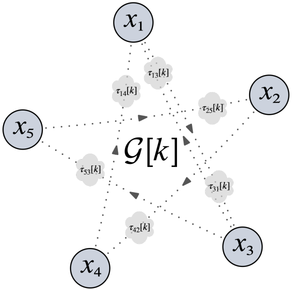

This often

- induces conservatism (discretization, input-delay)
- scales badly (# of decision variables may scale with ν)
- does not exploit the spatio-temporal interplay of the problem.

$$u_i(t) = K \sum_{j \in \mathcal{N}_i[k]} (x_i(s_k) - x_j(s_k))$$

A key observation: spatio-temporal interplay

Inter-neighbor interactions and local actions occur at different time scales.

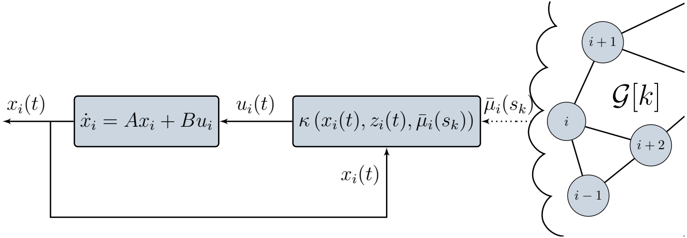

A key observation: spatio-temporal interplay

Inter-neighbor interactions and local actions occur at different time scales.

- Local information: the state $x_i(t)$ and controller variable $z_i(t)$, are continuously available.
  - Not effected by communication constraints!

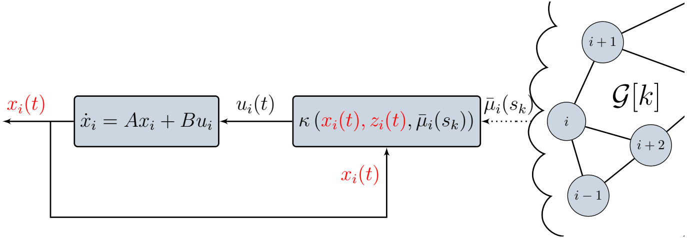

A key observation: spatio-temporal interplay

Inter-neighbor interactions and local actions occur at different time scales.

- Local information: the state $x_i(t)$ and controller variable $z_i(t)$, are continuously available.
  - Not effected by communication constraints!
- Neighboring information: transmitted information, $\bar{\mu}_j(s_k)$, is available intermittently and may be delayed.

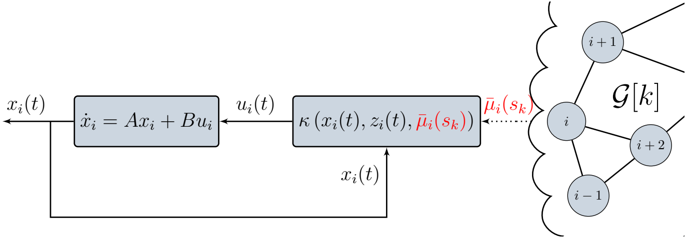

A key observation: spatio-temporal interplay

Inter-neighbor interactions and local actions occur at different time scales.

- Local information: the state $x_i(t)$ and controller variable $z_i(t)$, are continuously available.
  - Not effected by communication constraints!
- Neighboring information: transmitted information, $\bar{\mu}_j(s_k)$, is available intermittently and may be delayed.

A hybrid controller not not based on a discretized consensus protocol can exploit the interplay between local and global information.

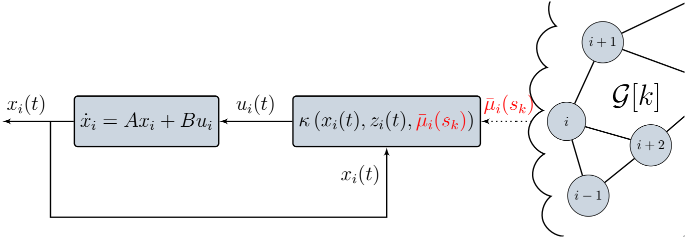

The paradigm: two measurements for two goals

- Flow dynamics: use local information to stabilize the system around a local *emulation* of the required trajectory $\bar{\mu}_i(t)$ via a continuous control signal.

The paradigm: two measurements for two goals

- Flow dynamics: use local information to stabilize the system around a local *emulation* of the required trajectory $\bar{\mu}_i(t)$ via a continuous control signal.

- Jump dynamics: at sampling instances, use incoming information, $\bar{\mu}_j(s_k), \in \mathcal{N}_i[k]$, to update the emulators discretely.

The paradigm: two measurements for two goals

- Flow dynamics: use local information to stabilize the system around a local *emulation* of the required trajectory $\bar{\mu}_i(t)$ via a continuous control signal.

$$\|x_i(t) - \bar{\mu}_i(t)\| \to 0, \quad \forall i$$

- Jump dynamics: at sampling instances, use incoming information, $\bar{\mu}_j(s_k), \in \mathcal{N}_i[k]$, to update the emulators discretely.

$$\|\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)\| \to 0, \ \forall i, j$$

- Two objectives: flow dynamics should track, and jump dynamics should synchronize.
  - This was the idea behind (Barkai, Mirkin, and Zelazo, 2023).

The hybrid controller for the delay-free problem

Theorem (Barkai, Mirkin, and Zelazo, 2023)

If $\mathcal{A}_{\mathbf{1,2}}$ holds, $A + BF_d$ is Hurwitz, and $A_0 = A + B\bar{F}$; then local control law

$$u_i(t) = F_d x_i(t) + \frac{1}{\nu}(\bar{F} - F_d)(z_i(t) + x_i(t))$$

generated by a hybrid controller with the following flow and jump dynamics

$$\begin{cases} \dot{z}_i(t) = (A + B\bar{F})z_i(t) + B(\bar{F}x_i(t) - u_i(t)), \quad z_i(0) = z_{i,0} \\ z_i(s_k^+) = z_i(s_k) - \frac{1}{\nu}\sum_{l \in \mathcal{N}_i[k]}(z_i(s_k) - z_l(s_k) + x_i(s_k) - x_l(s_k)) \end{cases} \tag{1}$$

asymptotically synchronize the agents to the required trajectory.

- Here the emulator is a linear combination of the agent and controller states

$$\bar{\mu}_i \coloneqq \frac{1}{\nu}(z_i + x_i)$$

Transmission delays in the hybrids setup

- The delays affect only transmitted information $\implies$ they do not affect the flow.

Transmission delays in the hybrids setup

- The delays affect only transmitted information $\implies$ they do not affect the flow.
- It can be shown that each emulator evolves according to

$$\begin{cases} \dot{\bar{\mu}}_i(t) = A_0 \bar{\mu}_i(t), \quad \bar{\mu}_i(0) = \bar{\mu}_{i,0} \\ \bar{\mu}_i(t_{ij}[k]^+) = \bar{\mu}_i(t_{ij}[k]) - \frac{1}{\gamma} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(t_{ij}[k]) - \bar{\mu}_j(s_k)) \end{cases}$$

equivalently, the jump dynamics can be written as

$$\bar{\mu}_i(t_{ij}[k]^+) = \mathrm{e}^{A_0 \tau_{ij}[k]} \left(\bar{\mu}_i(s_k) - \frac{1}{\gamma} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(s_k) - \mathrm{e}^{-A_0 \tau_{ij}[k]} \bar{\mu}_j(s_k)) \right)$$

Transmission delays in the hybrids setup

- The delays affect only transmitted information $\implies$ they do not affect the flow.
- It can be shown that each emulator evolves according to

$$\begin{cases} \dot{\bar{\mu}}_i(t) = A_0 \bar{\mu}_i(t), \quad \bar{\mu}_i(0) = \bar{\mu}_{i,0} \\ \bar{\mu}_i(t_{ij}[k]^+) = \bar{\mu}_i(t_{ij}[k]) - \frac{1}{\gamma} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(t_{ij}[k]) - \bar{\mu}_j(s_k)) \end{cases}$$

equivalently, the jump dynamics can be written as

$$\bar{\mu}_i(t_{ij}[k]^+) = \mathrm{e}^{A_0 \tau_{ij}[k]} \left(\bar{\mu}_i(s_k) - \frac{1}{\gamma} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(s_k) - \mathrm{e}^{-A_0 \tau_{ij}[k]} \bar{\mu}_j(s_k)) \right)$$

Transmission delays in the hybrids setup

- The delays affect only transmitted information $\implies$ they do not affect the flow.
- It can be shown that each emulator evolves according to

$$\begin{cases} \dot{\bar{\mu}}_i(t) = A_0 \bar{\mu}_i(t), \quad \bar{\mu}_i(0) = \bar{\mu}_{i,0} \\ \bar{\mu}_i(t_{ij}[k]^+) = \bar{\mu}_i(t_{ij}[k]) - \dfrac{1}{\gamma} \displaystyle\sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(t_{ij}[k]) - \bar{\mu}_j(s_k)) \end{cases}$$

equivalently, the jump dynamics can be written as

$$\bar{\mu}_i(t_{ij}[k]^+) = \mathrm{e}^{A_0 \tau_{ij}[k]} \left(\bar{\mu}_i(s_k) - \frac{1}{\gamma} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(s_k) - \mathrm{e}^{-A_0 \tau_{ij}[k]} \bar{\mu}_j(s_k)) \right)$$

We can exploit this to construct a simple predictor regardless of the gains $\bar{F}$ and F_d.

The main result

Theorem

If $\mathcal{A}_{\mathbf{1,2,3}}$ holds, $A + BF_d$ is Hurwitz, and $A_0 = A + B\bar{F}$; then the sampled-data controller (1) *with jump map*

$$z_i(t_{ij}[k]^+) = z_i(t_{ij}[k]) - \mathrm{e}^{A_0 \tau_{ij}[k]} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)),$$

asymptotically synchronize the agents for:

- *all initial conditions;*
- *all sampling sequences,*
- *and all time-varying delays satisfying $\mathcal{A}_{\mathbf{3}}$.*

Moreover, the system synchronizes with the same trajectory as in the delay-free case.

- Recall that $\bar{\mu}_i(t) = \frac{1}{\nu}(z_i(t) + x_i(t))$.

Proof outline - the big picture

1. It can be shown that $x_i(t)$ synchronize if and only if $\bar{\mu}_i(t)$ synchronize, and that

$$\begin{cases} \dot{\bar{\mu}}_i(t) = A_0\bar{\mu}_i(t), \quad \bar{\mu}_i(0) = \bar{\mu}_{i,0} \\ \bar{\mu}_i(t_{ij}[k]^+) = \bar{\mu}_i(t_{ij}[k]) - \dfrac{1}{\nu}\mathrm{e}^{A_0\tau_{ij}[k]}\displaystyle\sum_{j\in\mathcal{N}_i[t_{ij}[k]]}(\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)) \end{cases}.$$

   - Recall that $A_0 = A + B\bar{F}$.

Proof outline - the big picture

1. It can be shown that $x_i(t)$ synchronize if and only if $\bar{\mu}_i(t)$ synchronize, and that

$$\begin{cases} \dot{\bar{\mu}}_i(t) = A_0 \bar{\mu}_i(t), \quad \bar{\mu}_i(0) = \bar{\mu}_{i,0} \\ \bar{\mu}_i(t_{ij}[k]^+) = \bar{\mu}_i(t_{ij}[k]) - \frac{1}{\nu} \mathrm{e}^{A_0 \tau_{ij}[k]} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)) \end{cases}.$$

   - Recall that $A_0 = A + B\bar{F}$.

2. Equivalently, the new jump map can be written as

$$\bar{\mu}_i(t_{ij}[k]^+) = \mathrm{e}^{A_0 \tau_{ij}[k]} \left(\bar{\mu}_i(s_k) - \frac{1}{\nu} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)) \right),$$

which looks almost exactly like the delay-free case.

Proof outline - the prediction

1. $\mathcal{A}_3$ ensures that each agent can receive at most 1 update from each agent in each sampling interval.

Proof outline - the prediction

1. $\mathcal{A}_3$ ensures that each agent can receive at most 1 update from each agent in each sampling interval.
2. Assume that there are p delayed updates, and consider the sequence

$$s_k \leq q_1[k] < \cdots < q_p[k] < s_{k+1}$$

   where $q_i[k]$ are the ordered version of $t_{ij}[k]$.

Proof outline - the prediction

1. $\mathcal{A}_3$ ensures that each agent can receive at most 1 update from each agent in each sampling interval.
2. Assume that there are p delayed updates, and consider the sequence

$$s_k \leq q_1[k] < \cdots < q_p[k] < s_{k+1}$$

where $q_i[k]$ are the ordered version of $t_{ij}[k]$.

3. It can be shown that at the last update of each interval we have

$$\begin{aligned}\bar{\mu}_i(q_p[k]^+) &= \mathrm{e}^{A_0 \tau_p[k]} \left(\bar{\mu}_i(s_k) - \frac{1}{\nu} \sum_{l=1}^{p} \sum_{j \in \mathcal{N}_i[q_l[k]]} (\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)) \right) \\ &= \mathrm{e}^{A_0 \tau_p[k]} \left(\bar{\mu}_i(s_k) - \frac{1}{\nu} \sum_{j \in \mathcal{N}_i^{DF}[k]} (\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)) \right)\end{aligned}$$

where $\mathcal{N}_i^{DF}[k]$ is the neighborhood set of the delay-free system at $t = s_k$.

Simulation's setup

1. All Simulations involve $\nu = 3$ agent with the nominal sampling sequence randomized between 3 graphs and satisfy $\mathcal{A}_{\mathbf{2}}$.

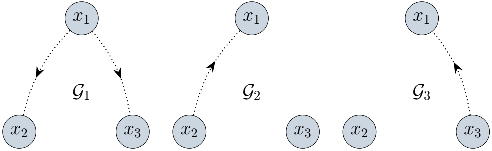

Simulation's setup

1. All Simulations involve $\nu = 3$ agent with the nominal sampling sequence randomized between 3 graphs and satisfy $\mathcal{A}_{\mathbf{2}}$.

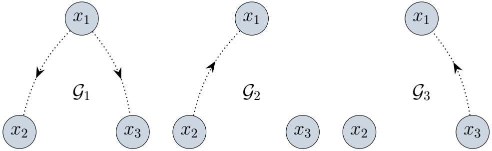

2. The nominal sampling intervals are randomized $s_{k+1} - s_k \in [0.3, 1.8]$, and the delays are generated by $\tau(s_{k+1} - s_k)$, where τ is uniformly distributed random variable from the interval $[0, 0.7]$, hence satisfy $\mathcal{A}_{\mathbf{3}}$.

Synchronization of LTI agents

Unstable LTI agents w/ time-varying trajectory

$$\dot{x}_i(t) = \begin{bmatrix} 4 & 9 \\ 1 & 4 \end{bmatrix} x_i(t) + \begin{bmatrix} 2 \\ 1 \end{bmatrix} u_i(t), \quad A_0 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

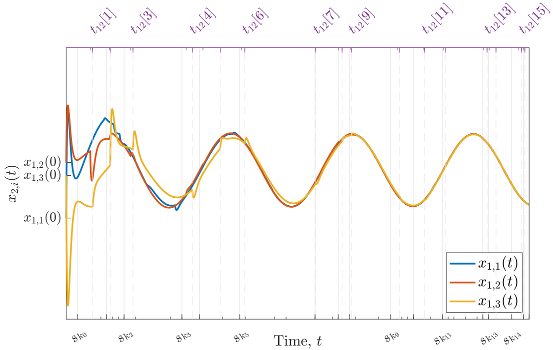

State's trajectories for first coordinate

Synchronization of LTI agents

Unstable LTI agents w/ time-varying trajectory

$$\dot{x}_i(t) = \begin{bmatrix} 4 & 9 \\ 1 & 4 \end{bmatrix} x_i(t) + \begin{bmatrix} 2 \\ 1 \end{bmatrix} u_i(t), \quad A_0 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

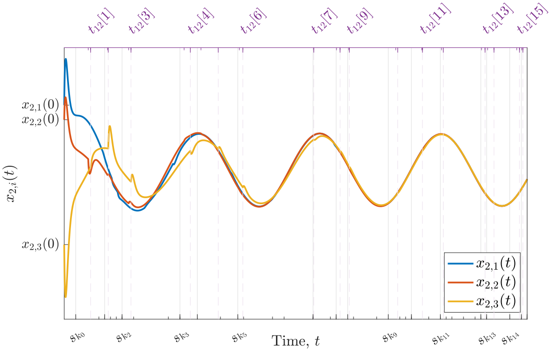

State's trajectories for second coordinate

Synchronization of LTI agents

Unstable LTI agents w/ time-varying trajectory

$$\dot{x}_i(t) = \begin{bmatrix} 4 & 9 \\ 1 & 4 \end{bmatrix} x_i(t) + \begin{bmatrix} 2 \\ 1 \end{bmatrix} u_i(t), \quad A_0 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

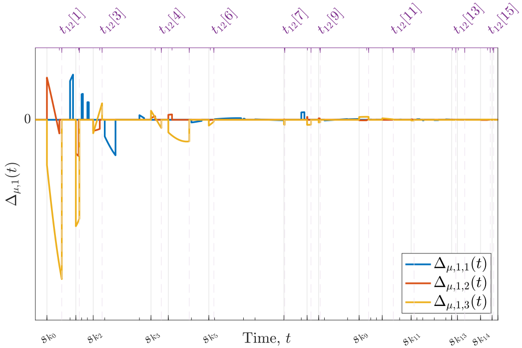

Error w/ delay-free system: $\Delta_{\mu,1i}(t) \coloneqq \bar{\mu}_{1i}(t) - \bar{\mu}_{1i}^{DF}(t)$

Synchronization of LTI agents

Unstable LTI agents w/ time-varying trajectory

$$\dot{x}_i(t) = \begin{bmatrix} 4 & 9 \\ 1 & 4 \end{bmatrix} x_i(t) + \begin{bmatrix} 2 \\ 1 \end{bmatrix} u_i(t), \quad A_0 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

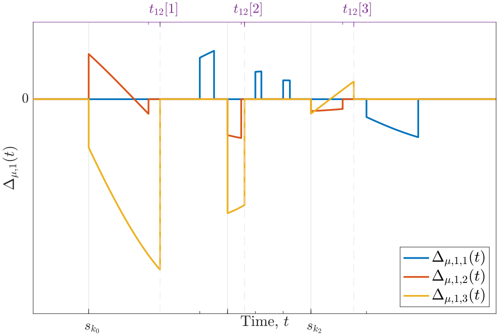

Error w/ delay-free system (zoomed): $\Delta_{\mu,1i}(t) \coloneqq \bar{\mu}_{1i}(t) - \bar{\mu}_{1i}^{DF}(t)$

Synchronization of LTI agents

Unstable LTI agents w/ time-varying trajectory

$$\dot{x}_i(t) = \begin{bmatrix} 4 & 9 \\ 1 & 4 \end{bmatrix} x_i(t) + \begin{bmatrix} 2 \\ 1 \end{bmatrix} u_i(t), \quad A_0 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

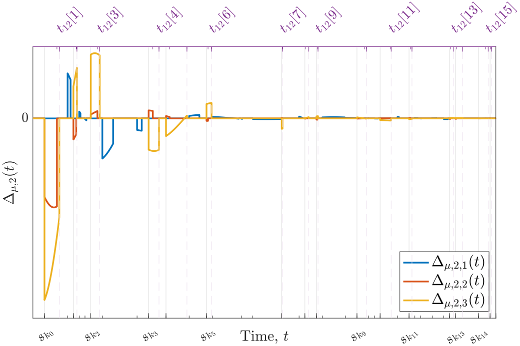

Error w/ delay-free system: $\Delta_{\mu,2i}(t) \coloneqq \bar{\mu}_{2i}(t) - \bar{\mu}_{2i}^{DF}(t)$

Concluding remarks

- Exploits the hybrid structure of the controller to compensate for transmission delays.
- Naturally works with heterogeneous and time-varying delays.
- The controller recovers the delay-free behavior - no loss of performance.
- It can be shown that each agent needs only a buffer of size 1 to implement the predictor.
- Future research:
  - Robustness to uncertainty in $\tau_{ij}[k]$
  - Extensions to longer time-delays
  - Moving from state to output feedback.